

#15

PATENT

REMARKS

Claims 17-39 were rejected under 35 U.S.C. §103(a) as being unpatentable over Gupta et al. U.S. Published Patent Application No. 2003/0196164 A1, in view of Shobatake et al., U.S. Patent No. 6,654,607. In view of the request for continued examination, and in an attempt to advance this lengthy prosecution to a close, the undersigned has reviewed the pending independent claims and made several new amendments for the Examiner's consideration. These amendments are not dictated by the prior art, as the pending rejection still appears misplaced given the actual subject matter claimed.

As previously noted, Gupta et al. describe a system whereby authors and end users can annotate media or other content delivered over the Internet. In Gupta et al., the content being annotated typically is provided by the streaming media server 11, and a separate annotation server 10 controls the delivery of the annotations, which are stored in an annotation content store 17. There is only one annotation server. A second data store 18 stores "meta data" associated with a given annotation. The meta data typically is in the form shown in Figure 5 of the application and includes, e.g., information identifying the author of the annotation, timestamps identifying to which portions of the media stream the annotation applies, the creation time, information identifying the annotation or related annotations, information identifying which versions of the media stream the annotation applies, and so forth. As illustrated in Figure 11, once a given media stream is being received at a client machine 15, the end user can obtain a previously-created annotation stored in the annotation content store 17. To this end, the annotation server identifies the media characteristics of the media stream, fetches the annotation, and serves that annotation to the requesting end user client machine. During this process, information (e.g., timestamp data) in the annotation meta data store is accessed and used by the annotation server to ensure that the annotation is synchronized with the media stream to which it applies.

Because the "annotation server" is the only device in Gupta et al. that applies information from the meta data store 18, the Examiner has equated the "annotation server" 10 with the described "content servers" of the "content delivery network." (See, e.g.,

#15

PATENT

claim 17). Moreover, the Examiner also appears to be equating the “annotation” itself with the “given content to be delivered over the content delivery network” and, further, that the information obtained from the meta data store 18 is the described “content control requirement.” These contentions are misplaced.

Respectfully, the annotation system described in Gupta et al. is not a “content delivery network” comprising (in amended claim 17, 25 and 33) “a plurality of content servers,” which content servers (in the plural) are used “on behalf of participating content providers” who have identified “given content to be delivered over the content delivery network.” (See, e.g., claim 17, 25 and 33). Indeed, as is evident from Figure 1, in Gupta et al. the content itself (typically a media stream) is obtained from the streaming media server 11 (and possibly the web page server 12), but there is one, and only one, “annotation server 10” for serving the annotations. Thus, even if one were to fairly construe an annotation author as a “content provider” and an “annotation” as “given content,” which Applicants do not concede, the claims as now amended require the invention to be implemented within the context of a “content delivery network” having “a plurality of content servers.” There is no disclosure or suggestion in Gupta et al. to use a distributed set of annotation servers. In this regard, the Examiner is directed to the limitation in each of independent claims 17 and 33 requiring that the “metadata for the given piece of content” be communicated “to the plurality of content servers.” Gupta et al. have no need for this step because they have just one annotation server, not a distributed set of such servers. Because Gupta et al. do not have multiple annotation servers and no need to communicate annotation meta data anywhere, they also fail to meet the limitations in dependent claims 19 (communication by header), 20 (communication by configuration file), 21 (communication by configuration file provisioned via an extranet application), 26 (communication), 27 (communication by a configuration file), 32 or 34 (communication by one of: a request string, a header and a configuration file), or 35 (communication by one of: a request string, a header and a configuration file, as provisioned via an extranet application).

The Examiner will also note that the “content delivery network” is now defined in amended claims 17 and 25 to comprise both the plurality of content servers and a domain

#15

PATENT

name service. (Independent claim 33 had this limitation originally). The Examiner's argument that Gupta et al. disclose a content delivery network because the reference discloses "delivering and rendering multimedia content in a network environment" (page 9, lines 1-2) is misplaced, respectfully. A "content delivery network" in the context of the present invention is a network managed by a service provider on behalf of participating content providers who desire to offload (to the CDN) certain content delivery requirements. As now amended, a content delivery network comprises the plurality of content servers and a domain name service that is authoritative for "domain[s] managed by the content delivery network service provider" (and, accordingly, is distinct from the Internet's basic DNS). Amended claims 17, 25 and 33 further require that DNS queries to the content provider domain are resolved by this content delivery network domain name service, which (as described later in the claim) does so using the domain managed by the CDN service provider in lieu (instead) of using the content provider domain. As noted above, Gupta et al.'s single annotation server is not a "plurality of content servers" and the reference does not teach having a CDN service provider with its own domain name service, let alone a DNS for resolving a "domain managed by the [CDN] service provider ... in lieu of the content provider domain." There are no such constructs or operations remotely suggested in Gupta et al.

Claims 17 and 25 have been further amended to reflect that a given content control requirement is a requirement to be applied to the given piece of content when the given piece of content is served from the content delivery network. (Claim 33 had this limitation originally). Gupta et al. do not teach a "content delivery network" so information in the annotation meta data store does not meet this limitation either.

Each of independent claims 17 and 25 now further require that the "given content control requirement specified in the metadata" be applied "at the given content server" conditionally. In particular, representative claim 17 describes "receiving a request for the given piece of content, determining whether a participating content provider has specified a content control requirement for the given piece of content and, if so, applying the content control requirement (specified in the metadata) prior to serving the given piece of content. Support for these limitations is provided, for example, in the discussion concerning Figure

#15

PATENT

5 on pages 32-33 of the written description. Gupta et al. have no concept of conditional application of an annotation (or annotation data) served from the annotation server. The conditional control feature now positively recited is advantageous as it enables participating content providers (which may be unaffiliated entities) to establish and set their own content handling requirements vis-à-vis the plurality of content servers (which, typically, are shared) on a per object basis for the content offloaded to the CDN.

Further, even if one were to fairly construe the Gupta et al. annotation as the “given piece of content,” the actual limitation requires that the “given content control requirement specified in the metadata” be applied at a content server that has been selected or identified from the “plurality of content servers” in the content delivery network. This requirement cannot be met in Gupta et al., as there is only one annotation server and, thus, no selection or identification of a “given” one selected or identified from a plurality of such servers.

As previously argued, it is also doubtful that one of ordinary skill in the art would read “given content control requirement” to reach the information that is stored in the annotation meta data store 18. In particular, the written description describes the CDN server metadata as a “set of control options and parameters that determine how a CDN content server will handle a request for an object” (see, e.g., text at page 8, line 21; page 2, lines 26-31 through page 3, line 2, emphasis supplied). Examples of such content control requirement metadata are provided, for example, on page 8, lines 25-30 and throughout the written description. For the reasons set forth above, the lone annotation server 10 is not a content delivery network (CDN) content server, i.e., one of a plurality of content servers to which participating content providers offload their content for delivery to requesting end users.

Turning to the dependent claim rejections, Gupta et al. does not teach communicating metadata in a configuration file. As per claim 21, an email is not a “configuration file” and there is nothing in Gupta et al. that suggests using an “extranet application,” such as a secure content provider customer portal, to generate a configuration file, which is then communicated to “a plurality of content servers.”

Claim 22 refers to applying an authentication-based or access control-based “content control requirement.” Thus, even if the annotation server is construed as the

#15

PATENT

“content server” and the “annotation” as the “given content,” there is nothing in Gupta et al. (especially the annotation handling routine of Figure 11) that indicates some authentication or access control method applied at the annotation server itself.

Regarding claims 23-24, the cited references do not distinguish in any way the meta data as “request” or “response” meta data.

The Examiner also is correct in concluding that Gupta et al. do not explicitly teach the steps of “having a participating content provider associate a content provider domain or subdomain with a domain managed by a content delivery network service provider” or “resolving a client query to the content provider domain or subdomain to an IP address of a given content server in the set of content servers using the domain managed by the content delivery network service provider.” Nevertheless, these limitations are said to be provided⁸ by the secondary reference, Shobatake et al. This contention remains incorrect, for the following reasons.

The Examiner contends that Shobatake et al. is analogous art because it “relates to the unification of various mobile communication protocols.” (See, comments on page 7). This technical description may be true, but it is not evidence that the reference is in the field of applicants’ endeavor (content delivery over the Internet) or reasonably pertinent to the problem sought to be solved (how to address content control requirements in an Internet-based content delivery network). (See, MPEP §2141.01(a)). As previously noted, the patent simply has nothing to do with Internet content delivery generally, content delivery networks in particular, or methods of controlling how content is served from content servers. In contrast, the patent describes how mobile devices (e.g., cellphones, PDAs, etc.) can be used across multiple, disparate communications platforms. In this system, mobile device users desire “the ability to communicate with anyone, irrespective of the network in which the called party (or callee is located). Because each of the mobile device networks have different protocols and technologies, no direct connection is possible.” (Column 2, lines 42+). Shobatake et al. address this problem by providing a unified mobility manager mobility that enables a mobile device to register with its “home database” even as the device moves within and across foreign networks.

#15

PATENT

The undersigned respectfully requests that the Examiner reconsider the conclusion that Shobatake et al. is analogous art that can support an obviousness rejection. It cannot. The Office has the burden of establishing a *prima facie* case of obviousness in the first instance. In view of the evidence of record – namely, that Shobatake et al. is classified in class 455 that is wholly unrelated to Internet or Web-based enabling technologies and infrastructure – it is respectfully submitted that the Office has not met this initial burden. Indeed, one of ordinary skill in the art of “content delivery” would not look to the mobile device telecommunications art to address the problems addressed by the present invention. As previously noted, the present invention deals generally with the problem of how content providers who are using a “content delivery network” can set up and ensure that given content control requirements are applied to the content being served by the CDN service provider’s content servers. This is a completely different problem from that addressed by Shobatake et al., who are addressing the problem of how to recognize a mobile device user as that user moves across disparate telecommunications networks.

Of greater importance is the lack of relevant subject matter in Shobatake et al. given the claimed invention. Shobatake et al. do not disclose the specific claimed requirements (see, e.g., claim 17) of “associating a content provider domain or subdomain with a domain managed by a content delivery network service provider” or using that domain to resolve something else, namely, a “DNS query to the content provider domain or subdomain.” In other words, these claims require the resolution of the “domain managed by [the] content delivery network service provider” where the DNS query is directed to some other domain or subdomain. This concept is nowhere taught or suggested by Shobatake et al, who merely disclose a set of alias databases 305-308; these databases, however, merely store “identification information relating to mobile terminals on other platforms.” (Column 5, lines 34-36). Thus, e.g., alias database 305 might store a cellular telephone user’s alias so that the mobility manager can recognize the mobile device user even as he or she moves into some foreign network. The claim limitations are not related to storing alias information per se; rather, as noted above, the claims require a specific type of alias (a “domain managed by a content delivery network service provider”) the processing of that alias (in lieu of some other Internet domain or subdomain actually

#15

PATENT

associated with the DNS query), as well as the further requirement that the domain processing identify a given content server from a "plurality of content servers" in a content delivery network.

Shobatake et al. cannot be combined with Gupta et al. absent some express teaching, suggestion or motivation to select the teachings of the separate references and combine them to produce the claimed combination. Once again, this is the Office's burden to show. With all due respect, the Examiner's continuing contention is this regard (obvious to combine "because user terminals querying a content provider domain with a domain managed by a content delivery network service provider ... would enhance the capabilities of Gupta by allowing for handling of communications in a unified manner") is unsupported and incorrect. In the first instance, as noted above, the claims do not say that "user terminals" query a "content provider domain with a domain managed" by the CDN service provider, as the Examiner states. Rather, the claims require that a given DNS query is to "the content provider domain or subdomain" and that such query be handled using not the "content provider domain or subdomain" in the query but, rather, the "domain managed" by the CDN service provider. Moreover, the "user terminals" in Shobatake et al. are mobile devices, such as cellphones or PDAs, and the problem of mobility management does not have anything to do with how Gupta et al.'s annotation system might be modified to derive the subject matter (as a whole) of any pending claim.

For these reasons, reconsideration and withdrawal of the obviousness rejection is requested.

Dependent claims 19-20, 26-27, 31-32 and 34 have been amended to conform the language therein to the changes to the parent claims.

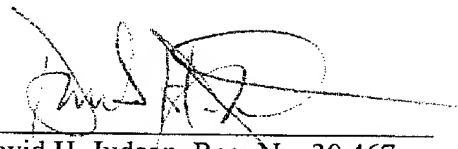
#15

PATENT

A Notice of Allowance is requested.

Respectfully submitted,

By:


David H. Judson, Reg. No. 30,467

15950 Dallas Parkway
Suite 225
Dallas, Texas 75248
(972) 385-2018